

A Framework for Targeting Banner Advertising On the Internet

Katherine Gallagher and Jeffrey Parsons

Faculty of Business Administration, Memorial University of Newfoundland

St. John's, NF, Canada A1B 3X5

{kgallagher, jeffreyp}@morgan.ucs.mun.ca

Abstract

Constraints that limit accurate targeting of advertising in traditional media may not hold in cyberspace. This paper presents a model for effectively and efficiently targeting hypermedia-based banner advertisements in an online information service. The model takes advantage of information technology to micro-target banner advertisements based on individual characteristics of users. A simple version of the model, which has the virtue of ease of development, is presented. Enhancements are also proposed. These require more effort to develop, but may lead to even more precise targeting of advertisements. Implementation of this framework may benefit both online advertisers and online consumers.

1. Introduction

Cyberspace is a rapidly growing new medium for commerce. To date, a great deal of industry attention has focused on electronic transactions over the Internet. Although rapid growth is predicted over the next few years [10, 17, 21], actual sales thus far have been only moderate: users appear to regard the Internet primarily as a source of product information--when it comes time to pay, they prefer to buy offline by more conventional means [12, 14].

Responding to consumers' desire for information, businesses in large numbers have developed sites on the World Wide Web (WWW or Web). Most commercial Web sites describe the firm and its products and/or services, and many offer opportunities for visitors to the Web site to provide feedback and ask for specific information. As well, some Web sites collect information from visitors in order to improve future offerings. Some sites also support ordering and payment. The interactive potential of Web sites is particularly exciting, as it facilitates relationship marketing and customer support, eliminating the obstacles of geography and time [14, 22]. Not surprisingly, then, industry and scholarly research has recently focused on making Web sites more appealing and useful to visitors [13]. However, a Web site can only be effective if current and prospective customers visit it. Attracting this audience is currently a major challenge.

In this paper, we address the challenge of attracting a defined target audience to a Web site via *banner advertising*. We propose a framework for effectively targeting banner advertising in an electronic marketplace in a manner that benefits both advertisers and consumers. It allows advertisers to reach consumers who are more likely to be interested in the products and/or services offered by the company, and exposes consumers to information about products and services that they are likely to be interested in purchasing. Although the framework is discussed in terms of the Internet, we believe it will be relevant to whatever form the "information superhighway" eventually assumes. The framework takes advantage of the capabilities afforded by information technology for collecting and processing data about users. The next section examines trends in the electronic marketplace. Subsequently, the current state of advertising in this medium is discussed. Thereafter, a framework for targeting banner advertising, supported by appropriate information technologies, is proposed. Finally, opportunities for further research are discussed.

2. Marketing and Advertising in an Evolving Electronic Marketplace

The Internet began in the early 1970s as a US government research project designed primarily for the needs of the military. It expanded in the 1980s to serve the international academic and research communities [19, 23]. In the 1990s, businesses began to appear on the Internet. Although accurate estimates are obsolete as soon as they are made, it is clear that today tens of millions of people have access to the Internet [16] through over 100,000 computer networks in 150 countries--and the numbers continue to increase [14]. Two types of developments are particularly noteworthy with regard to this growth.

First, a large and ever expanding number of affluent, educated consumers are using the Internet [11]. This concentration of very desirable consumers has led to a surge in commercial interest. Prior to 1990, nodes on the Internet were predominantly academic institutions. In 1990, about 1,000 businesses had Internet connections. By June 1995, over 21,000 businesses were online, and the growth in commercial connectivity shows no sign of slowing [8].

Second, the emergence of the hypermedia-based WWW,

together with point-and-click multimedia interfaces such as Netscape, have greatly increased usability of the Internet for persons without extensive computer training. The development of "applet" technology, such as Java, which allows programs to run on a variety of platforms, increases the transparency of various Internet services. In other words, as technology continues to evolve, it is no longer an obstacle to, but an enabler of, electronic commerce.

In this environment, companies are seeking ways to use the Internet effectively [1, 3, 13, 22]. One active area in electronic commerce involves using the Internet as a medium to communicate persuasive product and service information via advertisements. These take various forms, the most common of which are corporate Web sites and banner advertising. We define a banner advertisement as:

- " paid communication (via text, graphics, video and/or audio) of information about an organization and/or its products and services
- " by an identified sponsor
- " embedded within, and visually distinct from, information provided by an online service
- " with hypermedia links to the sponsor's Web site.

We distinguish banner advertising from simple hypermedia links (paid or not) to commercial Web sites: banner advertising conveys a message even if the user does not follow the link; simple links can only convey a message if the user follows the link. Banner advertisements are also distinct from what [14] refer to as "flat ads," single page advertisements that do not contain hypermedia links. In this paper, we restrict our discussion to banner advertising that appears in the course of users' browsing and searching activities on information services, such as Yahoo! (<http://www.yahoo.com>) and Excite (<http://www.excite.com>), that provide an entry point to Internet resources. Appendix 1 shows a banner advertisements by the Saturn automobile company.

Scant attention has been paid to banner advertising by researchers. This may be because banners seem relatively insignificant, especially when compared with the interactive richness of Web sites. Technical specifications for banner advertisements severely limit creative options and preclude any consumer-firm interaction beyond the consumer's selection of the hypermedia link to the associated Web site (Excite, for instance, specifies that "all banners are 468x60 pixels, gif format only, maximum file size is 7k" [9]). Banner advertisements are, however, very important and interesting when viewed as part of a *system* that converts browsers and searchers into Web site visitors and, ultimately, customers. In their model of this conversion process, Berthon, Pitt and Watson [3] identify a sequence of tasks. First, users must be made *aware* of the Web site, then they must be *attracted to* and *locate* the site. Once users have found the Web site, the task is to turn that hit into a

visit, ensuring there is some meaningful *contact* between the firm and the consumer; then to *convert* the visit into a purchase. The final task is to get purchasers to return to the Web site and *repurchase*. Each task in the sequence is dependent on the successful execution of the previous task.

Our view of the role of banner advertising in this system is as a mechanism to make target audience members aware of a firm's Web site and to attract those users to the site. We define two concepts critical to understanding this role. *Attraction effectiveness* is the number of target audience members who reach a company's Web site via a banner advertisement hypermedia link divided by the number of target audience members who use the information service on which the advertisement appears. *Attraction efficiency* is the advertising cost per target audience member attracted to a company's Web site via a banner advertisement.

There is some evidence that the attraction efficiency of banner advertising is low. A recent estimate indicates that only 1-2% of banner advertisements lead viewers to seek additional information (e.g., by selecting a hypermedia link to the company's Web site) [5]. Since information services charge advertisers based on number of exposures (e.g., [9, 24]), the cost of attracting a single target audience member to a Web site is at least 50 to 100 times what it would be if all users who were exposed to the advertisement selected the hypermedia link. (The cost is even higher if some users selecting the link are not target audience members.) Increasing attraction efficiency by reducing wasted exposures should therefore be a priority. (An additional motivation for improving performance of banner advertising in converting searchers and browsers into Web site visitors arises from recent events such as the agreement between Procter & Gamble and Yahoo! which provides for payment based on the number of people who actually seek additional information (by selecting a link from a banner advertisement) rather than those who are merely exposed to the advertisement [20]. Such arrangements are expected to pressure online services to eliminate wasted exposures [5].)

The estimate cited above does not provide evidence on the attraction effectiveness of banner advertising. The fact that only 1-2% of exposed users select a link to the advertiser's site is irrelevant to effectiveness if all target audience members using the information service are among this group. However, since banner advertisements on online information services are shown selectively to users, there will generally be the possibility that some target audience members who use the information service will not be exposed to the advertisement and, hence, will be unable to link to a company's Web site via it. Depending on the strategy used to select advertisements for users, a large number of target audience members may be missed.

We contend that both the attraction effectiveness and efficiency of banner advertising can be improved by

precisely *targeting* advertisements based on characteristics and behavior of individual users of information services. Moreover, such targeting can be more precise than the targeting possible in traditional media. For example, visitors to a "Travel" page on an information service may be good targets for an advertisement for discount airfares, as would readers of the Travel section of a newspaper. But the fact that the online visitors have made a series of decisions and taken a series of actions (i.e., selecting only a subset of highlighted links within a hierarchical menu of categories) to reach the Travel page, rather than some other page (e.g., the Home Decorating page) suggests they may have a greater interest in travel than, say, readers who unintentionally come upon the Travel section of a newspaper and decide to read it. Since these exposures are more likely to be target audience members, attraction effectiveness can be improved. Targeting individual users strategy should also lead to fewer wasted exposures, since the advertisement would not be shown to users who have not reached the Travel page, thereby improving attraction effectiveness. (See Appendix 2 for a similar example.)

At present, targeting of banner advertising does not always occur. For example, Appendix 3 shows an advertisement for Honda that appeared when Organic Gardening was selected from a hierarchical menu of categories. People interested in organic gardening may not be the best prospects for automobiles, as they are likely to be more environmentally sensitive than the general population and may feel that cars unnecessarily harm the environment.

Nevertheless, online information services do currently provide some targeting capability. As of August 1996, both Yahoo! [24] and Excite [9] offered advertisers three options: general rotation, geographic or content targeting, and keyword-based targeting. With "general rotation," banner advertisements rotate randomly through user searches and browsing on the site. The Honda advertisement that appeared on the Organic Gardening page in Appendix 3 was probably in general rotation. Restricted rotations permit advertisers to purchase space in specified content areas or by geographic region. For example, financial institutions can limit the exposure of their banner advertisements to users searching or browsing Business categories, and Canadian advertisers can choose to have their banner advertisements shown only to users who are searching or browsing in the Yahoo! Canada site. These two options are analogous to the targeting offered by traditional media such as newspapers, magazines, television, and radio [4].

The third option, keyword-based targeting, makes greater use of the targeting potential of information services. A company can buy keywords so that whenever a user enters one of those keywords during a search, s/he will be exposed to the company's banner advertisement. This ensures that the banner advertisement is presented only to people with a

demonstrated interest in the area. For instance, a marketer of golf equipment might buy the keyword "golf." Every time a user enters "golf" in a search, a banner advertisement for the equipment would appear. This is analogous to the more precise targeting provided by magazines.

While these are useful strategies, they fail to take full advantage of the targeting potential of banner advertising. Current technology provides the capability to develop sophisticated and detailed profiles of individual users of information services based on individual characteristics and past patterns of behavior in using the information service. The next section proposes and describes informally two versions of a model for targeting banner advertising by using the information technology on which an online information service is built.

3. A Model for Targeted Advertising

In traditional media, the quality of the information available constrains an advertiser's ability to target advertising effectively and efficiently. For example, many media buying decisions are based on data provided by research bureaus such as the Audit Bureau of Circulations (ABC), Business Publication Audit of Circulation (BPA), Arbitron, and A.C. Nielsen, which collect data on the demographics and media habits of consumers, and sometimes on product usage and brands [4]. These survey data are cross-tabulated to develop a profile of the audience of each media vehicle. The audience profile is then compared to the target audience profile identified by the advertiser to determine where there is a good match. For instance, an automobile manufacturer might identify the target audience for an advertisement for a particular model of car as middle-income females, 18 to 34, with busy lifestyles. Based on research bureau data, as well as the experience and judgement of the media planner, media vehicles with good reach in that demographic group would be chosen. Realistically, though, this type of targeting is usually very approximate. For instance, no matter how well the media vehicle audience profile matches the target audience profile, it is likely that only a portion of the audience would be in the market for a new car.

Online banner advertising may be able to overcome this problem. It is possible to target users very precisely because data can remain associated with individuals, so advertisers can select exactly the users to whom they wish their advertising to be exposed. It may be possible, for example, to identify which users will be in the market for a new car in a particular year. The remainder of this section describes two versions of a model for targeting banner advertising by taking advantage of the technological capabilities of the online environment. The model is designed to be appropriate for use by information services which sell

advertising space.

3.1. Basic version

The basic version of the model requires that users be assigned unique identifiers (e.g., user accounts) when they first connect to the information service. Subsequently, they provide these identifiers each time they connect. Users also complete an online questionnaire the first time they use the information service. (Incentives to complete the questionnaire may be provided by informing users that the information will be used to filter out advertising for products in which they are likely not to be interested.) The questionnaire allows data to be collected on several dimensions, including: (1) demographic attributes such as geographic location, income, family lifecycle stage, occupation, and sex; (2) psychographic attributes such as travel patterns and hobbies; and (3) product and brand usage attributes. This element of the basic model permits a banner advertisement to be directed to users (and only those users) who fit certain criteria, assuming data were collected on relevant attributes. For instance, a banner advertisement for baby strollers could reach parents of children under five years old--and only individuals in that group.

In contrast, research bureau data uses demographic correlates (e.g., males and females, 18 to 34) to identify media vehicles that attract a relatively large proportion of the people in the identified demographic group [4]. The media vehicles thus chosen may miss members of the target group (e.g., older parents) and reach consumers not in the target group (e.g., people who are between 18 and 34 but do not have young children). Even audience data based on cross-tabulations, while they supply information on more variables, still cannot isolate individuals who are in the target audience. (For example, research bureau data may allow an advertiser to identify a magazine whose audience includes a large number of people between 18 and 34 who have young children, but there will still be some readers who are not in the target market.)

The second element of the basic model involves eliciting the target audience profile from advertisers. An advertiser can specify a target audience using any number of attributes about which data have been collected. These can be expressed conjunctively and/or disjunctively. For example, a specification may indicate that an advertisement is to be presented to all users who (1) have household incomes over \$50,000, and (2) either work in a job that involves travel at least four times per year or have travelled on vacation in at least four of the past five years.

In this version of our model, the questionnaire determines the data collected about each user. The content of the questionnaire will vary depending on the nature of the information service, expected users, and expected

advertisers. However, it is imperative to design the instrument carefully, in consultation with advertisers based on anticipated relevant target audience attributes.

The final element of the model consists of a mechanism to select banner advertisements to display to users. The target audience profiles supplied by advertisers provide a screening mechanism over users. Each time a user connects, his/her profile is compared to all target audience profiles from all advertisers. The user's profile will actually match some subset of those profiles. If the number of matches is small (and the session is long), it will be feasible to display all banner advertisements associated with the matched profiles during the user's session. However, if the number of matches is larger (or the session is short), presenting all advertisements associated with the matched profiles may overwhelm the user. In such a case, it will be necessary to present only a selection of the identified target advertisements. A rationing system would be needed so that users are not deluged with banner advertisements while advertisers are assured of access to users who match the target audience profile.

In summary, the basic model has three elements: individual user profiles, individual advertisement target audience profiles, and a selection mechanism for presenting advertisements to specific users who match the target audience profile. This framework potentially eliminates wasted exposures and provides the capability to reach every single user who matches the target audience profile (this may not be realized if a rationing system is used). Users also benefit, since they will see advertisements only for products likely to be of interest to them.

3.2. Enhanced Version

The basic version of the model relies on users completing a questionnaire when they initially use an information service. This is a straightforward mechanism to collect data about user characteristics for the purpose of targeting advertisements. A similar approach has been incorporated in a commercial product for use with online catalogs to direct shoppers to products in which they are interested [15]. However, the advantage of simplicity is offset by several potential limitations. First, such information may become outdated, sometimes quickly, as user preferences and characteristics change. To some extent, information can be kept up-to-date by either readministering the questionnaire periodically or giving the user the opportunity to update her/his information (e.g., by a menu option or hypertext link) each time s/he connects to the information service. However, each of these strategies is intrusive and may impose an unwarranted burden on users in order to maintain currency of information.

A second, and perhaps more serious limitation of the

questionnaire strategy is that it is subject to two potential types of bias. First, the questionnaire designer will want to identify as many user attributes relevant to potential advertisers as possible. As the number of attributes increases, so does the length of the questionnaire, creating the possibility of higher mortality in completing the questionnaire (especially since it may be more difficult to induce users to complete it because they are both physically and psychologically remote), thereby increasing the potential nonresponse bias [7]. Second, the questionnaire method is plagued with well-known problems, such as errors due to inaccurate recall, telescoping, social desirability concerns, and cognitive biases, as well as ambiguity, intimidation, confusion, and incomprehensibility [2].

In view of these potential problems, it is appropriate to enhance the model so that it does not rely on user self-reports, can accommodate changing user characteristics and preferences, and is less constrained by the choice of questions. Fortunately, information technology may provide assistance in each of these areas.

Current technology allows a considerable amount of data about user search activities (both deliberate search and browsing) to be collected unobtrusively and analyzed to determine patterns. (We are dealing here only with the capabilities of the technology, not with the ethical issues such capabilities raise. However, we recognize that ethical issues must be considered explicitly in the design of systems based on our model. For instance, we believe users should be aware that such information may be collected, and how it may be used, and consent to this activity before using an information service.) In the enhanced model, we propose that patterns of search and browsing behavior exhibited by users while using an information service determine which advertisements are shown to that user during current or future sessions. In the remainder of this section, we provide a general overview of this approach.

As before, this model relies on assigning a unique identifier to each user for recording her/his searching and browsing activities while using the information service. Each session constitutes a "record", consisting of data such as: sites visited in order; pattern of navigation through a hierarchical category structure (as in Yahoo!); choice of search terms in keyword-based searches; and reaction to previously exposed targeted banner advertisements (e.g., which linked Web sites are selected and visited by the user and which ones ignored). The aggregate of such records for each user provides a profile from which preferences can be implicitly generated. As a simple example, if a user has made several searches using keywords such as "Atlantic salmon" and "fly fishing", and has visited the site of the *Angling Club Lax-a of Iceland* (http://www.ismennt.is/fyr_stofn/lax-a/uk/angl_uk.html), s/he may be targeted for a banner advertisement for a fishing

lodge in Alaska. However, if a user has previously been exposed to the same or similar banner advertisements but has not visited linked Web sites when there was an opportunity to do so, s/he may not be shown these banner advertisements in future.

This version of the model has the advantage of transparency. A user simply visits a service for whatever purpose s/he has in mind. Data are collected unobtrusively in the course of the visit. Moreover, the data reflect actual user behavior, rather than attitudes, intentions, or reported behavior captured through a questionnaire. Hence, the quality of data derived from user behavior should be superior to that of questionnaire data, for purposes of targeting advertisements.

A disadvantage of this model is the preparatory work involved on two fronts. First, it is not clear how to structure the data collected during visits so that useful information can easily be coded for storage and later extraction. Research is needed to develop useful and efficient coding mechanisms for storing such data as sequences of visits and search terms used. We expect this can be handled using conventional database structures such as relations (tables); however, the design of a relational database for this purpose is itself a distinct research issue. Second, the ability to store the required data does not necessarily mean useful information can be extracted from it. Further research is required to determine the types of analyses that yield insights into user characteristics and preferences hidden in the data.

The enhanced model should be used in conjunction with the basic model. A questionnaire may be very effective for identifying various demographic data relevant to advertisers but impossible to ascertain simply from users' online search and browsing behavior. However, since demographic data has limitations for effectively targeting consumers of most products, the enhanced model of data collection may yield complementary data on preferences from patterns of online search and browsing behavior.

The next section describes an implementation architecture for the basic version of the model. Extensions that support the enhanced version of the model remain as future research.

4. An Implementation Architecture

The architecture required to implement the basic version of the model consists of two parts: data structure to represent user profiles and target audience profiles, and an algorithm to select banner advertisements to display to a user. This section describes these components.



Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.